09-6-05 AF

Atty. Docket No. UND005
Client Matter No. 83208.0008
Express Mail Label No. EV544475953US

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

In re Application of:	Confirmation No. 5465
Gregory Hagan Moulton, et al.	Art Unit: 2151
Serial No.: 09/777,002	Examiner: Hassan Phillips
Filed: February 5, 2001	Customer No. **25235**
For: **SYSTEM AND METHOD FOR INTELLIGENT, GLOBALLY DISTRIBUTED NETWORK STORAGE**	Docket No. UND005

TRANSMITTAL OF BRIEF ON APPEAL UNDER 37 C.F.R. § 41.37

Mail Stop Appeal Brief – Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Please find enclosed an Appellants' Brief for the above application. Also, please find enclosed a check in the amount of $250.00 for a small entity. Any fee deficiency associated with this submittal may be charged to Deposit Account No. 50-1123.

Respectfully submitted,

September 2, 2005

Kent A. Lembke
Atty. Reg. No. 44,866
HOGAN & HARTSON LLP
One Tabor Center
1200 17th Street, Suite 1500
Denver, Colorado 80202
(720) 406-5378 Tel
(720) 406-5301 Fax

Attorney Docket No. UND005
Client/Matter No. 83208.0008
Express Mail No. EV544475953US

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

In re Application of:	Confirmation No. 5465
Gregory Hagan Moulton, et al.	Art Unit: 2151
Serial No.: 09/777,002	Examiner: Hassan Phillips
Filed: February 5, 2001	Customer No. **25235**
For: **SYSTEM AND METHOD FOR INTELLIGENT, GLOBALLY DISTRIBUTED NETWORK STORAGE**	Docket No. UND005

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPELLANTS' BRIEF UNDER 37 CFR § 41.37

I. Real Party in Interest

Avamar Technologies, Inc.
135 Technology Drive, Suite 100
Irvine, CA 92618
USA

II. Related Appeals and Interferences

No other appeals or interferences are currently known to Appellants that will directly affect, be directly affected by, or have a bearing on the decision to be rendered by the Board of Patent Appeals and Interferences in the present appeal.

III. Status of Claims

Claims 1-7 and 12-30 are pending in the application. No claims have been allowed with claims 1-7 and 12-30 being rejected.

Claims 8-11 and 31 are cancelled.

The rejection of claims 1-7 and 12-30 is being appealed.

IV. Status of Amendments

A claim amendment canceling previously withdrawn claim 31 was submitted after the final Office Action mailed April 21, 2005. In the Advisory Action mailed July 6, 2005, the Examiner indicated that for purposes of appeal that the amendment would be entered. As a result, Appellants believe that all claim amendments have been entered.

Claims 1-7 and 12-30 as pending are provided in the Claims Appendix.

V. Summary of Claimed Subject Matter

Claims 1, 14, and 27 are independent claims that are being appealed.

Initially, Appellants believe it would be useful to understand the problems addressed by the invention claimed in independent claims 1, 14, and 27. As noted in Appellant's Background at line 13, page 4 to line 33, page 5, the invention addresses problems associated with the inherent limitations of RAID and NAS storage that make it difficult to effectively and strategically locate data in storage devices in a distributed storage network. "The strategic location of data storage refers to the process of determining a location or locations for data storage that provide a specified degree of availability, reliability, and security based on the relevant contexts associated with the data storage facilities" (emphasis added). The claimed invention addresses the need for data storage systems and methods that take into account storage criteria in a user's data storage request to select storage nodes or devices based on the "context" of the particular storage nodes. Examples of such node selection based on node or device contexts that together satisfy the request criteria.

With specific reference to claim 1, a data storage system is specified that includes a "plurality of storage nodes, each node existing at a physical location and having one or more associated contexts including a political context, an economic context, a geographic context or a network topological context." With reference to Appellants' specification, the storage nodes are represented by nodes 101, 102, 103, 104, and 105 of Figure 1, and each of such nodes may be provided by the devices shown in Figure 3, e.g., by RAIN elements 315. As discussed in the first paragraph of the Detailed Description (page 8, lines 7-25), storage nodes and/or capacity are associated with "contexts" and the "contexts are represented and communicated by the dynamic exchange of state information between the storage nodes" such that desired criteria (e.g., "requirements placed on the data associated with one or more set of desired criteria") can be fluidly met or satisfied by storage of user data in one or more storage nodes having context "to maintain dynamic compliance with the set of desired criteria."

The "associated contexts" or "current context of the storage nodes" may be "represented by state information held in the state information data structure" (see, page 21, lines 27-29 of Appellants' specification). Figure 6 illustrates a state information data structure or system state information element 603 that contains the context associated with each storage node that is provided by state information, which in this exemplary representation includes access speed, transfer rate, network locality, physical locality, inter-connectedness, security, reliability, political domain, capacity, cost, and other attributes.

The "associated contexts" of claim 1 include "a political context, an economic context, a geographic context or a network topological context." A general definition of each these contexts is provided with the description of state information or other attribute information that may make up each of these particular storage node contexts in Appellants' application from line 19, page 4 to line 19, page 5.

Claim 1 also calls for the storage system to include "data storage management processes that select one or more of the storage nodes to serve a data storage request based at least in part upon the particular contexts associated with each of the storage nodes." Figure 5 illustrates an iRAIN component 502 that may be used to process

data storage requests from client 501 to select storage nodes (in RAIN 505) based on the contexts associated with nodes and the storage request. Additionally, Figure 6 illustrates the iRAIN component 502 with task processing 602 that interacts with system state information 603 to determine which nodes should be used to satisfy the storage request based on the criteria of the request and the contexts of the nodes. The processing of storage requests is further explained in Appellants' specification at page 21, lines 8-29, with iRAIN processes 502 using the desired criteria in the storage request to find nodes that exist in contexts that match the criteria (with current context of storage nodes being represented by held state information for each node). For example, if the storage request specified a particular cost criteria, the data storage management processes would attempt to select a storage node with an "economic context" that matches the cost criteria in the storage request (as indicated by held node state information).

Independent claims 14 and 27 are directed to methods of managing distributed data storage system with limitations similar to that of claim 1 presented in method form. As a result, the summary of the claimed invention for claim 1 is believed applicable to claims 14 and 27, with each claim involving using contexts described by state information and associated with storage nodes to allocate capacity of the distributed storage nodes to satisfy a data storage request from a file system or client device.

VI. Grounds of Rejection to Reviewed on Appeal

1. Claims 1-7 and 13-20 stand rejected under 35 U.S.C. §103(a) as being unpatentable over U.S. Pat. No. 5,987,506 ("Carter").

2. Claim 12 stands rejected under 35 U.S.C. §103(a) as being unpatentable over Carter in view of U.S. Pat. No. 5,794,254 ("McClain").

VII. Argument

A. Rejection of Claims 1-7 and 13-30 Under 35 U.S.C. §103(a) Based on Carter is Improper

In the final Office Action mailed April 21, 2005, claims 1-7 and 13-30 were rejected under 35 U.S.C. §103(a) as being unpatentable over U.S. Pat. No. 5,987,506 ("Carter"). This rejection is traversed based on the following remarks, and Appellants request that the rejection be reversed as not properly supported by Carter.

Claim 1 is directed to a data storage system. Carter fails to show or suggest a storage system with "a plurality of storage nodes...having one or more associated contexts including a political context, an economical context, a geographical context or a network topological context." As explained previously, such "contexts" may be associated with a node by storage of state information as shown in Figure 6 in the state information structure or system 603. Further, Carter fails to show or make obvious a storage system with "data storage management processes that select one or more of the storage nodes to serve a data storage request based at least in part upon the particular contexts associated with each of the storage nodes" (emphasis added). The system of claim 1 utilizes the associated contexts to identify or select among a plurality of storage nodes "to serve a data storage request." Because at least these two elements are missing from the teaching of Carter, the rejection of claim 1 based on Carter is improper and should be withdrawn.

The April 21, 2005 Office Action states at page 4 that Carter teaches a data storage system with storage nodes that each have "one or more associated contexts" at col. 6, lines 7-12. Further, at page 2 in the Response to Arguments, the Examiner cites to Carter at col. 6, lines 10-17 for showing "contexts" with its "globally unique identifier" that is used to access a particular node. Carter is directed to a method of accessing a shared memory structure made of a number of nodes by a plurality of clients, but Carter fails to teach or suggest "contexts" that are associated with each storage node. Instead, Carter teaches "a globally addressable storage system that allows remote computers and computers on different, interconnected networks to communicate and share data" (see, col. 3, lines 1-4). To allow each storage node to be accessed, a unique identifier is assigned to each node that allows other devices to

find and/or access the node. Carter at col. 7, line 42 to col. 8, line 49 with reference to Figure 2 explains in detail the mapping of a directory and file structure to an addressable shared memory system made up of a number of storage devices. Study of Figure 7 along with Carter's specification beginning at col. 20, line 57 is also useful for understanding that this reference is directed to shared memory systems and managing access of the shared memory (e.g., with a directory manager 244 and a flow scheduler 272).

There is no teaching, however, at this portion of Carter or elsewhere that each node has associated with it "contexts" as called for in claim 1 and as construed consistently with this terms plain meaning and in line with Appellants' specification and figures. The Examiner is citing an identifier useful for addressing a storage device but "contexts" as called for in claim 1 and shown in Figure 6 are attributes or state information for a device that can be used in selecting the storage node to satisfy a storage request but not for addressing the node to store data. Further, as pointed out by Appellants in the Amendment of June 21, 2005, "the specific examples of context recited in claim 1 refer to externalities that define context in which the storage nodes exist. These examples of context do not refer to the data itself (e.g., a file name, a file system name, etc.) but instead represent external contextual information." In summary, there is no teaching in Carter of associating "contexts" with each node for later use in selecting that node based on a storage request.

Further, claim 1 calls for the "contexts" to specifically include "a political context, an economical context, a geographical context or a network topological context," and these specific contexts are not shown or suggested by Carter. In the Response to Arguments, the Office Action on page 3 indicates that "Carter does not show storage nodes having associated contexts including a political context, an economic context, a geographic context, or a network topological context." The Examiner, however, asserted that it would have been obvious to one of ordinary skill in "the field of politics, economics, geography, or network topology would wish to store data associated with politics, economics, geography, or network topology." Thus, the Examiner asserted that one skilled in the art could modify Carter to "use the globally unique identifier to access a particular node that will serve the data storage request."

Appellants assert that the Examiner's argument is faulty, and it would not be obvious to modify Carter to achieve the claimed invention. First, the Examiner states it would be obvious that people in certain fields of study or work would wish to store data related to that field. However, that is not what is called for in claim 1. Claim 1 calls for a context to be associated with each storage node and for such associated contexts to provide specific information on the node. For example, the associated context may be an economic context for the node or a political context for the node. This is not calling for storing economic or political data in the storage node but is instead providing an economic or political context for the storage node. The Examiner then states with such stored data that the globally unique identifier can be used to access the appropriate node to serve a data storage request. However, the identifier simply allows each storage node or device to be addressed within a shared memory structure, and the identifier does not provide context (such as political, economic, geographic, or network topological context) that can be used meet a storage request. For this additional reason, Carter does not support a rejection of claim 1.

In the Advisory Action, the Examiner states that it "is not clear that 'context' does not refer to the data itself," which he indicates enables him to take a broad interpretation of the claims to have Carter show the claimed contexts. Appellants disagree as claim 1 is written such that the contexts are "associated" with "each node" and not with the data stored on such nodes (with no antecedent basis provided for stored data in the first element of claim 1). Hence, this argument of an allowable broad construction of claim 1 is faulty and shows that the Examiner is incorrectly construing the claim language.

Further, on page 5 of the Office Action, the Examiner states "having the context include political, economic, geographic, or topological context is a field of use limitation, and not patentable distinction." The Appellants have requested clarification on this rejection in the past Amendments as follows: "it is respectfully requested that the Office supply sum [sic] legal precedent for this assertion. The patent laws do not include any provisions that support the position taken in the Office action. Moreover, claim 1 does not attempt to claim the existence of context, but instead is claiming the use of this context information in the service of a data storage

request. It is respectfully requested that this argument either be supported by statutory or case law, or be withdrawn." The Examiner only responded to this request by stating in the Advisory Action that "Applicant is directed to 35 USC 103(a)," which is silent on "field of use limitations." Hence, Appellants again request that the rejection of claim 1 be reversed as there is no support in Carter for contexts associated with each storage node and clearly, no support for the particular contexts listed in the claim.

Yet further, Claim 1 calls for "data storage management processes that select one or more of the storage nodes to serve a data storage request based at least in part upon the particular contexts associated with each of the storage nodes." The Office Action cites Carter at col. 7, lines 43-49 for providing teaching such management processes. However, at this citation, Carter discusses data control programs that couple with shared memory system "similarly to an interface between a conventional data storage program and a local memory device." A conventional storage program does not select one or more storage nodes based on context associated with each node. Further, the Carter system does not teach that its management mechanisms select nodes in its shared data system to serve a storage request based on context associated with the nodes. More clearly, Carter does not discuss that such associated node used in the selection process would be economic, geographic, political, or network topological. Apparently, Carter simply pools the storage devices and then maps requests to all devices. Hence, Carter fails to teach the data storage management processes of claim 1.

Claims 2-7 and 13 depend from claim 1 and are believed allowable over Carter at least for the reasons provided for allowing claim 1. Further, claim 5 calls for the data storage requests to be "associated with a set of desired criteria" and the data storage management processes include "processes for matching the desired criteria to the contexts of the storage nodes." Carter fails to teach such a matching of criteria in a storage request with contexts of storage nodes as part of a process of selecting which nodes to use to serve a data storage request. The Office Action cites Carter at col. 6, lines 7-14 for teaching the limitations of claim 5. However, Carter at this citation is merely discussing the use of accessing a node based on its unique identifier or address. There is no discussion of matching criteria in a storage request with

context associated with the node to select one or more nodes for use in serving the request but instead Carter simply discusses how a node can be accessed by a data control program (without discussing how each node is selected for use in storing data).

Dependent claims 6 and 7 are directed to selecting the nodes based on their contexts and presenting "a unitary logical volume of data storage" to the devices requesting data storage. The Office Action cites Carter at col. 7, lines 61-67 for teaching this feature, but Carter is at this citation is discussing using a graphical user interface to graphically depict the shared store of data (element 28 of Figure 1). There is no teaching of processes to present a unitary logical volume and that such a volume "satisfies the desired criteria" as called for in claim 7. Hence, Carter fails to support a rejection of these claims.

Independent claim 14 is directed to a method of managing distributed data storage and includes elements in method form that are similar to those of claim 1. As a result, Appellants believe that the reasons provided for allowing claim 1 over Carter are equally applicable to claim 14. Claim 14 includes "determining desired criteria associated with the received data storage task" and then selecting the storage nodes "having an associated context satisfying the desired criteria." As discussed with reference to claim 1, Carter teaches using a shared memory structure but fails to teach selecting nodes in the structure to satisfy a storage task and particularly, selecting such nodes based on criteria in the storage task and contexts associated with the nodes that meet these criteria. Carter fails to teach or suggest each and every element of claim 14, and hence, Carter does not support a rejection of claim 14 based on obviousness.

Claims 15-26 depend from claim 14 and are believed allowable as depending from an allowable base claim. Further, claim 15 requires that the selected nodes comprise two nodes that separately would not meet the criteria of the storage task but together do. Such a limitation cannot be met by Carter's unique identifier which the Examiner states is a "context", which only provides a way of accessing a device in a shared storage network. The Office Action cites Carter at col. 7, lines 8-38 for teaching this limitation, but at this citation, Carter is discussing using multiple devices

or nodes to make up a shared memory space (element 20 of Figure 1). There is no discussion here, though, that the Carter system looks at the contexts of two devices to find a combination that jointly satisfies the criteria of a storage task.

Dependent claim 22 calls for the selection of the nodes to be based on the "socio-economic attributes of the physical location of the data storage node." There is no teaching in Carter of such a selection technique as all the nodes in Carter are simply pooled for apparently arbitrary access. The Office Action cites Carter at col. 6, lines 13-14 in rejecting claim 22. Carter states at this citation "control program can access a desired portion of the structured store using a globally unique identifier." This provides absolutely no suggestion of using socio-economic attributes of the physical location of the storage node to select which among a plurality of nodes to use to satisfy criteria associated with a data storage task.

Independent claim 27 is directed to a data storage service method that includes limitations similar to that of claims 1 and 14, and the reasons provided for allowing claims 1 and 14 over Carter are believed applicable to claim 27. Particularly, claim 27 calls for contexts of a number of network-accessible storage devices to be maintained. The contexts are kept by maintaining "a state information data structure including state information" that describes these contexts. The state information is then used to allocate capacity within the storage devices "to handle the received data storage request."

As discussed with reference to claims 1 and 14, Carter fails to teach associating contexts for a number of storage devices and then using such contexts to select which devices are used to handle received data storage requests. The Office Action cites col. 13, lines 19-48 for teaching such a state information data structure, but at this citation, Carter only discusses a "lock object data structure" for which current states of lock are kept to allow a file system (element 60) to coordinate accessing of shared files. This teaching does not show maintaining state information for a number of storage devices to describe each of their contexts (including a political context, an economic context, a geographic context, or a network topological context). Instead the lock state information corresponds to files and ranges of files not to contexts of storage nodes that can be used in selecting storage nodes to handle a

received data storage request. Because Carter fails to show "maintaining a state information data structure..." or "using the state information to allocate capacity..." Carter fails to support a rejection of claim 27. Claims 28-30 depend from claim 27 and are believed allowable as depending from an allowable base claim.

Based on the above remarks, Appellants respectfully request that the rejection of claims 1-7 and 13-30 be reversed.

B. Rejection of Claim 12 Under 35 U.S.C. §103(a) Based on Carter and McClain is Improper

Also, in the Office Action dated April 21, 2005, claim 12 was rejected under 35 U.S.C. §103(a) as being unpatentable over Carter in view of U.S. Pat. No. 5,794,254 ("McClain"). Claim 12 depends from claim 1 and is believed allowable as depending from an allowable base claim. Further, McClain is cited for its teaching of encrypting messages and fails to overcome the deficiencies associated with Carter discussed in the reasons for allowing claim 1 over Carter. For these reasons, Appellants request that the rejection of claim 12 based on Carter and McClain be reversed.

Conclusion

In view of all of the above, claims 1-7 and 12-30 are believed to be allowable and the case in condition for allowance. Appellants respectfully request that the Examiner's rejections based on 35 U.S.C. §103 be reversed for all pending claims.

Respectfully submitted,

Date: September 1, 2005

Kent A. Lembke, Reg. No. 44,866
HOGAN & HARTSON LLP
One Tabor Center
1200 17th Street, Suite 1500
Denver, Colorado 80202
Phone: (720) 406-5378
Fax: (720) 406-5301

VIII. CLAIMS APPENDIX

1. A data storage system comprising:

a plurality of storage nodes, each node existing at a physical location and having one or more associated contexts including a political context, an economic context, a geographic context or a network topological context;

interface mechanisms coupled to each storage node for communicating storage access requests with the storage node; and

data storage management processes that select one or more of the storage nodes to serve a data storage request based at least in part upon the particular contexts associated with each of the storage nodes.

2. The system of claim 1 wherein the data storage management processes comprise computer-implemented processes executing in at least one of the storage nodes.

3. The system of claim 1 wherein the data storage management processes comprise computer-implemented processes executing in all of the storage nodes.

4. The system of claim 1 wherein the data storage requests are associated with a set of desired criteria.

5. The system of claim 1 wherein the data storage requests are associated with a set of desired criteria and the data storage management processes comprise processes for matching the desired criteria to the contexts of the storage nodes.

6. The data storage system of claim 1 wherein the data storage management processes present a unitary logical volume of data storage to external devices generating the storage access requests to the selected one or more storage nodes.

7. The data storage system of claim 6 wherein the selected one or more storage nodes are selected such that the contexts associated with the unitary logical volume satisfies the desired criteria associated with the one or more storage nodes.

Claims 8 - 11 (cancelled)

12. The data storage system of claim 1 further comprising encryption mechanisms coupled to the interface mechanisms for encrypting storage access requests during communication between nodes.

13. The data storage system of claim 1 further comprising authentication mechanisms coupled to the interface mechanisms for authenticating storage nodes before communicating storage requests.

14. A method of managing distributed data storage comprising the acts of:

providing a plurality of distributed storage nodes, each node existing at a physical location and having one or more associated contexts including a political context, an economic context, a geographic context or a network topological context;

receiving a data storage task in one of the storage nodes;

determining desired criteria associated with the received data storage task;

selecting one or more of the plurality of storage nodes having an associated context satisfying the desired criteria; and

executing the storage task in the one or more selected storage nodes.

15. The method of claim 14 wherein the selected storage nodes comprise at least two storage nodes where neither of the at least two storage nodes individually satisfy the desired criteria, but collectively the at least two storage nodes satisfy the desired criteria.

16. The method of claim 14 wherein the selected storage nodes comprise at least two storage nodes and the at least two storage nodes are located in different geographical locations.

17. The method of claim 14 wherein the selected storage nodes comprise at least two storage nodes and the at least two storage nodes are located in different areas of a single data center.

18. The method of claim 14 wherein the selected storage nodes comprise at least two storage nodes and the at least two storage nodes are connected via different network backbones in a single data center.

19. The method of claim 14 wherein the selected storage nodes comprise at least two storage nodes and the at least two storage nodes are located in different data centers.

20. The method of claim 14 wherein the selected storage nodes comprise at least two storage nodes and the at least two storage nodes are located in different cities.

21. The method of claim 14 wherein the selected storage nodes comprise at least two storage nodes and the at least two storage nodes are located in different political jurisdictions.

22. The method of claim 14 wherein the selection is based upon socio-economic attributes of the physical location of the data storage node.

23. The method of claim 14 wherein the selecting step further comprises matching the desired criteria to a context associated with the one or more storage nodes.

24. The method of claim 14 wherein the act of storing comprises storing the data according to a distributed parity scheme analogous to parity distribution found in RAID subsystems.

25. The method of claim 24 wherein the parity paradigm comprises an N-dimensional parity mechanisms where "N" is greater than three.

26. The method of claim 14 wherein the act of storing comprises storing the data in a manner such that the data stored in any one storage node cannot be used in any meaningful fashion without the availability of some or all of the data stored in other storage nodes.

27. A data storage service comprising:

receiving data storage access requests from a file system;

maintaining a state information data structure including state information describing the contexts of a number of network-accessible storage devices, wherein the contexts include a political context, an economic context, a geographic context or a network topological context; and

using the state information to allocate capacity within the network-accessible storage devices to handle the received data storage request.

28. The data storage service of claim 27 wherein the act of maintaining the state information data structure comprises:

detecting a change in state information associated with at least one of the network-accessible storage devices; and

updating the state information associated with the at least one network-accessible storage device to include the change in state information.

29. The data storage service of claim 27 further comprising:

dynamically re-allocating capacity within the network-accessible storage devices in response to detecting a change in their associated contexts.

30. The data storage service of claim 29 wherein the dynamic re-allocating is done in the absence of an externally generated data storage access request.

Claim 31 (cancelled)

IX. EVIDENCE APPENDIX

No copies of evidence are required with this Appeal Brief. Appellants have not relied upon any evidence submitted under 37 C.F.R. §§ 1.130, 1.131, or 1.132.

X. RELATED PROCEEDINGS APPENDIX

There are no copies of decisions rendered by a court or the Board to provide with this Appeal as there are no related proceedings.